

Greedy Compositional Clustering for Unsupervised Learning of Hierarchical Compositional Models

Adam Kortylewski, Clemens Blumer, Thomas Vetter

Department of Mathematics and Computer Science,
University of Basel

Abstract. This paper proposes to integrate a feature pursuit learning process into a greedy bottom-up learning scheme. The algorithm combines the benefits of bottom-up and top-down approaches for learning hierarchical models: It allows to induce the hierarchical structure of objects in an unsupervised manner, while avoiding a hard decision on the activation of parts. We follow the principle of compositionality by assembling higher-order parts from elements of lower layers in the hierarchy. The parts are learned greedily with an EM-type process that iterates between image encoding and part re-learning. The process stops when a candidate part is not able to find a free niche in the image. The algorithm proceeds layer by layer in a bottom-up manner until no further compositions are found. A subsequent top-down process composes the learned hierarchical shape vocabulary into a holistic object model. Experimental evaluation of the approach shows state-of-the-art performance on a domain adaptation task. Moreover, we demonstrate the capability of learning complex, semantically meaningful hierarchical compositional models without supervision.

Keywords: Unsupervised Learning, Hierarchical Compositional Models

1 Introduction

Compositionality is a fundamental principle underlying human visual pattern analysis. Objects are seamlessly decomposed into parts, constrained by spatial as well as hierarchical relationships. This process offers advantages for vision systems in terms of efficient representation of patterns as well as concerning the recognition of patterns based on top-down contextual constraints. A major challenge, however, is to infer the compositional dependence structure of patterns from images without supervision. A number of works have successfully demonstrated the ability to perform this learning task [1–4]. The learned models have shown impressive generalization abilities for a diverse set of applications such as image classification [1], object parsing [2], transfer learning [3] or one-shot learning [4]. We see two main learning paradigms underlying these diverse set of learning algorithms. Either the hierarchy is learned bottom-up, layer by layer, until no further compositions are found [2, 1] (Figure 1a), possibly followed by a top down pass to correct for small errors in the model [2]. Alternatively, a template for the whole object is learned directly and subsequently divided into parts according to a predefined




Fig. 1: Schematic illustration of different iterative procedures for learning hierarchical compositional models. Parts are organized in multiple layers and represented by colored boxes. Green: Hidden states are inferred; Gray: Fixed states; Blue: Complete distribution over node states available; White: Not involved in the current learning step. Red arrows depict learning processes, black arrows are pre-defined rules. (a) Bottom up structure induction as e.g. applied in [5, 2]. (b) Feature pursuit as proposed in [3]. (c) The proposed greedy EM-type algorithm applied for bottom-up compositional learning.

splitting scheme [3], thus generating the intermediate layers of the hierarchy (Figure 1b). Each of these learning paradigms however suffers from one of two fundamental drawbacks, that have severe implications for the learning procedure. The bottom-up strategy is constrained by the assumption that object parts have to be detected in order to learn the next higher order layer in the hierarchy (Figure 1a). However, local ambiguities in the data often lead to wrong or missing detections that affect the learning of all subsequent layers. Top-down knowledge from higher-order parts would be helpful for resolving these ambiguities. On the other hand, the top-down learning scheme as applied in [3] assumes that the structure of the model is known beforehand (Figure 1b), thus arguably the main challenge in unsupervised learning of hierarchical models is solved with supervision. Moreover, most patterns have different structural dependencies, rendering this approach unfeasible when the number of objects is high or even unknown.

We propose a greedy EM-type algorithm that combines the best of both learning schemes: A bottom-up induction of the compositional structure of patterns, while avoiding a hard decision on the activations of the hidden variables. Our approach builds on the active basis model that was originally introduced in [6, 7] and was successfully adjusted to handle unaligned data in [8, 9, 3]. We propose to integrate this learning framework into a greedy bottom-up structure induction process that will allow the induction of the compositional structures of patterns. Figure 1 schematically illustrates the main learning schemes. In Figure 2 the main differences to the active basis model framework are visualized. The main novelties of this work are:

1. A method that integrates a bottom-up dictionary learning with a top-down model building process into a joint hierarchical compositional learning framework.
2. A greedy unsupervised learning scheme that automatically infers the number of active basis models needed to explain a given set of images
3. A bottom-up structure induction procedure that directly works on the response maps of parts, and thus does not rely on hard detections.

This paper is organized as follows: Section 2 embeds the proposed work into the context of hierarchical models and deformable templates. Section 3 covers the learning of hi-

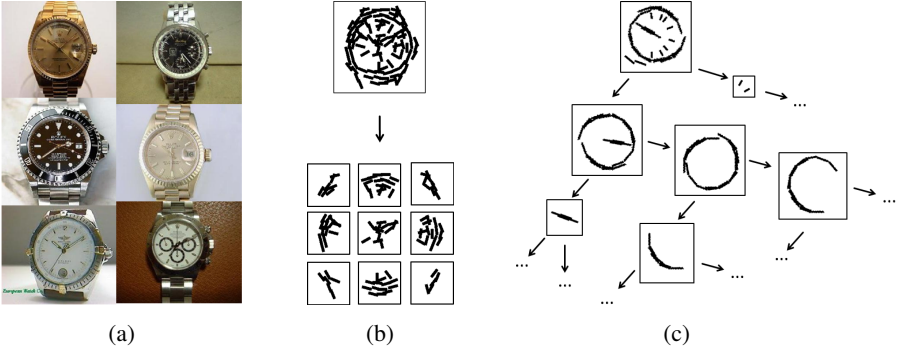


Fig. 2: Two hierarchical compositional models and their decompositions into individual parts. The models are learned from 10 watch images. (a) A sample of the training data; (b) The resulting model learned with the approach from [3]; (c) the result of the proposed approach. The Gabor wavelets are illustrated with a black bar at the corresponding location, scale and orientation. In contrast to the pre-fixed decomposition scheme in (b) the model in (c) is decomposed in a semantically more meaningful way.

erarchical compositional active basis models. Subsection 3.1 introduces the theoretical background on active basis models. Each of the subsections 3.2, 3.3 and 3.4 explains a learning step of the proposed greedy EM-type algorithm (also illustrated in Figure 1c). In Section 4 we demonstrate that the learned object representation achieves state-of-the-art performance on a challenging transfer learning task. Moreover, we are able to learn semantically meaningful hierarchical object decompositions without supervision. The code for the learning framework can be downloaded from [website].

2 Related Work

Hierarchical Models: Hierarchical object models have a long history in computer vision, at least dating back to Fukushima’s neocognitron [10]. In the last decade, hierarchical models have been successfully applied in a diverse set of computer vision applications, e.g. in [11–15]. However, these models were trained with detailed supervision or are based on hierarchical clustering of abstract features [16–18]. The principle of compositionality, which was e.g. extensively studied in [19], enabled an efficient unsupervised learning of hierarchical models. These models have been shown to achieve state-of-the-art performance for different computer vision tasks [16, 18, 13, 3]. In contrast to [13, 3, 9], where at least part of the structure of the models was manually set, we induce the whole structure of the model automatically. The work of Filder et al. [5, 1] and L. Zhu et al. [2] demonstrates that the structure of fully generative hierarchical compositional models can be learned with no supervision. However, both approaches suffer from the fact that in order to learn e.g. a layer at layer L the learned parts at layer $L - 1$ have to be detected first. Due to local ambiguities in the image wrong or missing detections can occur, thus disturbing the input to the learning for all subsequent layers. In contrast to that, our approach directly works on the response maps of the parts, thus

avoiding these early decisions, which results in a less error-prone structure induction process.

Deformable Templates: Deformable templates have been very successfully applied for analyzing objects under local perturbations e.g. in [20, 21]. Recently, a particular deformable template named ‘active basis model’ has been proposed within an elegant information-theoretic framework [7]. Inspired by Olshausen and Field’s linear sparse coding model [22] and local maximum pooling of feature responses suggested in the work of Riesenhuber and Poggio [23]. Object shapes are modeled as compositions of Gabor wavelets that can locally perturb their location and orientation in order to compensate for local shape variations. The learning framework is based on the matching pursuit algorithm proposed by Mallat and Zhang [24]. A hierarchical active basis model was initially mentioned in [8] and then proposed in [9, 3]. A main constraint of the approach is however, that the number of parts as well as the number of layers in the hierarchy have to be fixed beforehand. Thus, the main challenge in the process of unsupervised learning of compositional models is solved with supervision. Furthermore, it is assumed that the parts are fixed at canonical locations with respect to the object center (e.g. arranged in a 2×2 grid) and that these locations are the same for every object. In contrast to that, we propose to integrate the active basis model learning scheme into a bottom-up hierarchical learning framework. This permits the induction of the compositional model structure including the number of parts their spatial arrangement and the depth of the hierarchy.

3 Learning Hierarchical Compositional Models

Our goal is to learn generative hierarchical compositional models from images without supervision. The proposed algorithm integrates the active basis model framework [7] into a bottom-up learning scheme in the spirit of [5, 2]. In the following Section 3.1 we shall introduce the theoretical background for active basis models. More information on this is available in the work in [7, 6, 8]. Each of the following Sections explains one step of the proposed learning scheme in the same order as illustrated in Figure 1c. Starting from the original learning approach for single active basis models (Section 3.2), we present an extension that allows the automatic determination of the optimal number of active basis models in Section 3.3. In Section 3.4, we integrate the greedy principle into a bottom-up structure induction process. Section 3.5 introduces a top-down process that composes the dictionary of hierarchical templates into a holistic object representation.

3.1 Active Basis Models - Theoretical Background

Active basis models are a type of deformable templates for describing object shapes under small local shape deformations. The template is composed of basis filters in a certain global spatial configuration. Each filter, however, is active in the sense that its location and orientation can be varied independently of the other filters. Formally, it is a linear additive model in the form of the well-known sparse coding principle proposed

by Olshausen and Field [22]. It is an extension to the original work, in the sense that it is applied to represent an ensemble of images:

$$I_m = C_m B_m + U_m = \sum_{i=1}^N c_{i,m} B_{i,m} + U_m, \quad (1)$$

where $\{I_m, m = 1, \dots, M\}$ are image patches of size $d \times d$ which are reconstructed by a set of Gabor wavelets in certain positions and orientations $B_{i,m}$. Beware that $B_i = B_{X_i, s_i, \alpha_i, (0,1)}$ denotes a pair of even and odd ($\{0, 1\}$) Gabor wavelets at a certain position X_i , frequency $1/s_i$ and orientation α_i . The corresponding coefficients are $c_{i,m}$ and the residual image is denoted by U_m . The power of the model lies in the property that each filter is allowed to actively perturb its location and orientation within a specified range $\Delta = \{\delta X, \delta \alpha\}$, thus $B_{i,m} = \{B_{X_{i,m} + \delta X, \alpha_{i,m} + \delta \alpha}\}$. In this way, the model can compensate small object deformations without re-optimizing the states of all variables, as would be e.g. the case when using a global pattern model. Following the original algorithm, the wavelets have zero mean unit l_2 norm and a fixed central frequency. We assume the patches I_m are aligned and depict an object in a fixed pose. We will relax this assumption to unaligned objects in Section 3.2. A detailed theoretical underpinning and probabilistic formulation of active basis models is given in [7, 8]. In the following we summarize those results that are most important for the understanding of this work.

We model the probability of a patch I_m given a set of basis filters B_m as:

$$\begin{aligned} p(I_m|B_m) &= p(C_m, U_m|B_m) \\ &= p(U_m|C_m, B_m)p(C_m|B_m) \\ &= p(U_m|C_m, B_m) \prod_{i=1}^N p_i(c_{m,i}). \end{aligned} \quad (2)$$

The factorization in the last step is based on the assumption that the model has a compositional tree structure and that parts do not overlap. Let q be a background distribution modeling the distribution of coefficients and residual images as they occur in random natural images. We model the probability of an image I_m under q as $q(I_m) = q(C_m, U_m) = q(U_m|C_m)q(C_m)$. During model learning, we want to maximize the difference between the object model $p(I_m|B_m)$ and the background distribution $q(I_m)$. We assume that the probability distributions of the residual background are the same $q(U_m|C_m) = p(U_m|C_m)$ as in [7, 8]. Basically this means that those parts of the image that cannot be explained by the Gabor wavelets B_m follow the same distribution. The likelihood ratio between the object and the background then becomes:

$$\frac{p(I_m|B_m)}{q(I_m)} = \frac{p(U_m|C_m, B_m) \prod_{i=1}^N p_i(c_{m,i})}{q(U_m|C_m, B_m) \prod_{i=1}^N q(c_{m,i})} = \prod_{i=1}^N \frac{p_i(c_{m,i})}{q(c_{m,i})}. \quad (3)$$

The reference distribution $q(c)$ can be estimated by pooling a histogram of basis filter responses from a random set of natural images. We assume q is stationary and thus translation-, rotation- and scale-invariant. We estimate $p_i(c_{m,i}) = p(c_{m,i}, \lambda_i)$ by the

following exponential family model:

$$p(c_{m,i}, \lambda_i) = \frac{\exp(\lambda_i z(|c_{m,i}|^2) q(c_{m,i}))}{Z(\lambda_i)}, \quad (4)$$

As proposed in [7] we apply a sigmoid transform $z(r) = \tau[2/(1 + e^{-2r/\tau}) - 1]$ to $r = |\langle I_m, B_{i,m} \rangle|^2$ that saturates at value τ in order to discount for large values of r . The normalizing constant $Z(\lambda_i)$ can be estimated by integrating the numerator on a set of natural training images. Following the maximum entropy principle [25], the maximum likelihood estimate for the natural parameter is $\hat{\lambda} = \mu^{-1}(\sum_{m=1}^M z(r_{m,i})/M)$, where $\mu = \frac{1}{M} \sum_{m=1}^M z(r_{m,i})$ can be estimated from the patches I_m . Given a patch I , the log-likelihood ratio between the object model with parameters $\Theta = \{\lambda_i, B_i\}$ and the background distribution is:

$$l(I|\Theta) = \sum_{i=1}^N [\lambda_i \max_{\Delta} z(|\langle I, B_{X_i+\delta X, \alpha_i+\delta\alpha} \rangle|^2) - \log Z(\lambda_i)]. \quad (5)$$

3.2 Learning Dictionaries of Active Basis Models

In this Section we will explain a feature pursuit algorithm [7] to fit the model in Equation 1. The procedure corresponds to the first step in the illustration of our learning scheme from Figure 1c. In the literature [7, 6, 3], different names have been introduced in order to distinguish between algorithms that follow the same principle but differ in mathematical details. However, these details are not important for the understanding of this work. Therefore, we will confine ourselves to describing just the so-called shared sketch principle, which is illustrated in pseudo-code in Algorithm 1. More detailed information can be found in the original works.

Algorithm 1 Shared Sketch Principle

Input: Set of response maps $\{R_{1,1}^\Delta, \dots, R_{N,M}^\Delta\}$; Set of filters $\{B_1, \dots, B_N\}$;

E : Number of filters to be composed.

Output: A spatial arrangement of filters $a = B_1, \dots, B_E$ that sequentially maximize $r_e = \sum_{m=1}^M R_{e,m}^\Delta$.

- 1: **for** $e = 1, \dots, E$ **do**
 - 2: Select the filter that maximizes r_e .
 - 3: Suppression: Set $R_{i,m}^\Delta \leftarrow 0$, for all positions where $|B_e|^2 > \tau$
-

The main goal behind the shared sketch principle is to determine a spatial arrangement a of filters from the set B_m such that the filters in a sketch as many edges in the training data I_m as possible. The process is initialized with a set of response maps $S = \{R_1, \dots, R_{M,N}^\Delta\}$. A response map $R_m^\Delta = \max_{\Delta} z(|\langle I_m, B_{X_i+\delta X, \alpha_i+\delta\alpha} \rangle|^2)$ is computed by convolution of one image from I_m with a filter from B_m , subsequently followed by a maximization over the perturbations $\Delta = \{\delta X, \delta\alpha\}$ and the application

of the sigmoid transform z . The algorithm continues by sequentially selecting a basis filter B_i that maximizes the sum of transformed filter responses $r_i = \sum_{m=1}^M R_{i,m}^A$ over the set of images I_m . Subsequently, all overlapping filters are suppressed before choosing the next one. This procedure is repeated until a fixed number of filters has been chosen. In the previous Section we assumed the patches I_m to be aligned. The assumption can be relaxed by integrating the procedure into an EM-type learning scheme [8]. Let assume we want to learn a dictionary with a fixed number of K active basis models $A = \{a_1, \dots, a_K\}$. For each a_k , we perform the following steps in parallel until the learned models converge to a stable solution:

Algorithm 2 EM-type Dictionary Learning

- 1: Randomly sample M patches I_m^k from the image for each model $a_k \in A$.
 - 2: Learning: Learn the active basis models A with the shared matching pursuit process based on I_m^k (Algorithm 1).
 - 3: Encoding: Evaluate the log-likelihood ratio from Equation 5 based on the learned models A at each position in the images.
 - 4: Cut new patches I_m^k from those M positions with the highest score. Go back to step 2.
-

In practice, the model converges after a few iterations. In the next Section, we extend this EM-type learning scheme with a greedy selection scheme in order to automatically determine the number of models K . Furthermore, we will integrate the greedy learning into a bottom-up structure induction process in Section 3.4, thus learning dictionaries of compositional active basis models.

3.3 Greedy Dictionary Learning

The procedure described in Algorithm 2 is the basis for a mostly unsupervised way of learning active basis models. However, some supervision is still required in the sense that the number of models must be fixed beforehand. One way to determine the number of models is by evaluating the adjusted Bayesian information criterion [8] for multiple runs of the algorithm. However, this brute-force type of learning is time-consuming. Additionally, the random initialization makes the learning process highly stochastic, which leads to very different results for different trials. Thus, the resulting independence structure is highly sensitive to the initialization. Both issues are undesirable for unsupervised vision systems. In this Section, we propose a learning scheme that enables the automatic determination of the number of active basis models. In Figure 1c, this is depicted in the second step of the structure induction process. In order to overcome the issues, we propose to learn the models greedily one at a time, instead of learning them all at once in parallel (see function `GreedyDictionaryLearning` in Algorithm 3). We start with one active basis model and cut an initial set of image patches randomly from the training set as before. When the first model is converged, we divide the dictionary pool A into two disjoint sets $A = \{A_F, A_B\}$. The learned template is shifted to A_B and is from then on constrained to only take part in the encoding step, while the set A_F contains the template that is currently learned. In the encoding step,

A_B acts as a background process that explains away parts of the image. The model in A_F is more likely to explain different parts of the image with a higher score than A_B and is thus more likely to specialize into a new type of shape. The greedy learning is stopped as soon as the learned template in A_F is no longer able to encode any new parts of the image with a higher score than those in A_B . This greedy procedure also enables us to change the initialization from randomly sampling image patches to a more goal-oriented sampling scheme. One possible approach is e.g. to sample patches at positions where the log-likelihood ratio of the models in A_B is low.

In this Section, we explained one main novelty of the proposed work, the greedy dictionary learning scheme. It solves two main issues: The number of active basis models can be determined automatically and the learned independence structures are much more similar for different runs of the learning procedure. The presented greedy learning scheme is the basis for the compositional active basis model that we shall introduce in the following.

3.4 Bottom-up Compositional Learning of Hierarchical Models

In this Section we introduce a bottom-up learning scheme for compositional active basis models with a hierarchical matching pursuit process. We start by learning simple compositions of Gabor wavelets and recursively apply the proposed greedy learning scheme in order to learn higher order compositions. New parts are again allowed to perturb locally, thus generating active basis models composed of active basis models. The resulting dictionary of compositional active basis models can be further integrated into a holistic object model with a top-down process we shall present in Section 3.5. Given a set of training images $T = \{t_1, \dots, t_N\}$ and a dictionary of Gabor wavelets B , we want to learn a dictionary of hierarchical compositions of active basis models $H = \{h_1, \dots, h_L\}$. The layer h_1 is the set of all parts that are composed of Gabor wavelets. h_2 contains the compositions of elements from h_1 , et cetera. We will refer to $h_l = \{A_1^l, \dots, A_{N_l}^l\}$ as layers and to elements of h_l as parts. In order to simplify the notation, we assume the sets h_l also include rotated versions of the learned parts. Following the work in [3], we refer to the response maps of the models in h_l on a set of images in as SUMmaps R_{h_l} . In order to account for the active perturbation, we maximize the SUMmaps over the parameter ranges as specified in $\Delta^L = \{\delta_{X_L}, \delta_{\alpha_L}\}$ before the learning procedure. The resulting response maps $R_{h_l}^\Delta$ are referred to as MAXmap. In the following Section we shall explain the learning procedure in more detail. Algorithm 3 sketches the process in pseudo-code while Figure 3 provides a visual illustration.

Hierarchical Composition. The hierarchical structure induction process is initialized with a dictionary of active basis models $h_1 = \{a_1^1, \dots, a_{N_1}^1\}$ composed of Gabor wavelets B . This initial dictionary is learned with the greedy EM-type algorithm as described in Section 3.3. The set h_1 defines the first layer of the hierarchy (see Figure 3). The next higher layer composes elements of h_1 into higher order parts. Whereas h_1 was learned on the response maps of the Gabor wavelets R_B^Δ , h_2 is now learned on the response maps $R_{h_1}^\Delta$. In the beginning, patches are randomly cut from $R_{h_1}^\Delta$ and an initial active basis model a_1^2 is learned based on the shared sketch principle. This time however the set h_1 is used as basis filters. This initial model is iteratively refined with the introduced EM-type learning. After a certain number of iterations, a_1^2 is fixed and a

Algorithm 3 Greedy Hierarchical Compositional Learning**Input:** A set of layer one models h_1 and the corresponding response maps MAX_1 .**Output:** A Hierarchical Compositional Dictionary $H = \{h_2, \dots, h_L\}$.

```

1:  $L = 1$ ;
2: do
3:   Start learning next layer  $L = L + 1$ 
4:    $[h_L, MAX_L] \leftarrow \text{GreedyDictionaryLearning}(h_{L-1}, MAX_{L-1})$ 
5: while found composition  $h_L$ 

6: function GREEDYDICTIONARYLEARNING(  $h, R$  )
7:    $A_B = \{\}, A_F = \{\}$ 
8:   do
9:     Cut out an initial set of patches  $R_m$  from the response maps  $R$ 
10:     $A_F \leftarrow \text{sharedSketch}(R_m, h)$ 
11:    for #iterations do
12:      SUMmap  $\leftarrow$  Evaluate each model from  $\{A_F, A_B\}$  on  $R$ 
13:      MAXmap  $\leftarrow$  Maximize SUMmap over  $\Delta^L$ 
14:       $R_m \leftarrow$  cut  $M$  patches from maximal positions in MAXmap
15:       $A_F \leftarrow \text{sharedSketch}(R_m, h)$ 
16:      Add the learned model  $A_F$  to the set  $A_B$ 
17:   while found a composition  $A_F$ 

```

new model is learned on randomly cut patches. From this point on, a_1^2 is no longer re-learned. However, it still competes for image resources during the encoding step. The learning stops as soon as a new initial model is no longer able to gather enough patches during the encoding step. This implies that the model is highly redundant as it is not able to explain any part of the images better than the models learned already. The resulting dictionary of active basis models h_2 is subsequently used as a basis for learning the next higher layer. The procedure is recursively applied until no new compositions are found. We refer to this process as greedy hierarchical compositional learning (Algorithm 3).

Top-Down Information. An important feature of the proposed learning scheme is that it operates entirely on the response maps. Other approaches usually take an early decision whether a part $a_{N_l}^l$ is present in an image or not directly after the layer l is learned. This is typically done by detecting parts according to a pre-fixed response threshold. However, this often leads to wrong decisions. Due to local ambiguities, false-positive detections could occur. Even worse, parts that respond slightly below the threshold would be missed and thus cannot serve as evidence for higher-order parts. Those missing parts could be recovered with a top-down correction process as e.g. proposed in [2]. However, this is based on the assumption that the bottom-up process converges to a result that is reasonably close to the correct solution. In contrast to that, in the proposed learning scheme the complete response maps $R_{h_l}^\Delta$ are available during the learning of layer h_{l+1} . Thus, the critical decision on the presence of parts $a_{N_l}^l$ is avoided until top-down information from higher order parts $a_{N_{l+1}}^{l+1}$ is available.

Conclusion. In this Section we presented a bottom up compositional learning scheme that hierarchically composes local deformable templates into higher order deformable




Fig. 3: Illustration of the joint bottom-up and top-down compositional learning scheme. During the bottom-up process (blue box) higher order parts are built by composing parts from the same layer based on the proposed greedy learning algorithm (Algorithm 3). When no further compositions can be found, the top-down process is initialized (green box). The part from the highest layer 5 is greedily augmented with lower layer parts based on the same greedy learning algorithm.

templates. During the learning process we can avoid the hard detection of parts. This marks the second main novelty of the proposed work. The resulting dictionary of hierarchical active basis models is generative and fully probabilistic. In order to compose the individual dictionary elements into a hierarchical compositional model we propose to apply the greedy dictionary learning scheme in a top-down manner. In the next Section we shall explain this procedure in detail.

3.5 From Dictionaries to Object Models

So far, we have demonstrated the ability to learn individual hierarchical deformable templates. The task of constructing a full object model from part models is far from being solved yet, although it has been studied actively. Recent approaches in the context of HCMs can be found e.g. in [2, 1, 9]. In this section we will show that the same greedy learning principle we used during the bottom-up process, can also be applied in a top-down manner. The result of this procedure is a holistic object model composed of the individual parts from H . The main purpose of this section is to demonstrate the universal applicability of the proposed greedy learning algorithm for bottom-up as well as top-down learning. A more elaborate analysis is beyond the scope of this paper.

During the bottom-up compositional learning, different parts of the object terminate at different layers of the hierarchy. E.g. in Figure 3 the hour markings on the dial of the watch are represented at layer two. However, the circular shape of the watch is com-

posed of more elements and is therefore represented at a higher layer in the hierarchy. In order to learn a representation of the complete watch, these parts have to be composed into a holistic object model. The basic idea of the proposed top-down grouping process is to augment the top layer part with models from the layers below. We do so by greedily learning compositions according to the EM-type learning scheme as introduced in Section 3.3 until no further compositions are found. In order to make the top-down model building possible a few algorithmic adjustments have to be made compared to the bottom-up process. The algorithm starts at the top layer and then proceeds down the hierarchy. We start by choosing one part from the top layer as an initial seed to be augmented $s \in h_L$. The goal is to augment s with parts from the layers below $\{h_{L-1}, \dots, h_1\}$. Subsequently, we perform the greedy dictionary learning as proposed in Algorithm 3 with the set of basis filters being $\{s, h_{L-1}\}$. Additionally, we define that the first chosen filter in the shared sketch procedure must be s . The remaining $E - 1$ filters must be selected from h_{L-1} . In this way we induce an augmentation mechanism that composes s with elements from h_{L-1} . The top-down process proceeds with the next lower layer when no further compositions are found. Figure 3 illustrates the integration of the bottom-up and top-down processes, which is the third main contribution of this work. Although, many open questions remain regarding e.g. the modeling of articulated objects or whole object classes, our goal is to demonstrate that the proposed greedy EM-type algorithm naturally integrates bottom-up and top-down learning of pattern representations. In the next Section, we provide a thorough evaluation of the presented learning framework.

4 Experiments

In the following Section the proposed greedy EM-type algorithm will be evaluated based on two experiments. First, it will be applied to unsupervised learning of hierarchical compositional object representations from images. Subsequently, the algorithm is evaluated on the task of domain adaptation on the four domain dataset. In order to guarantee the reproducibility of the work, the code for the learning algorithm as well as for the experiments can be downloaded from [website]. Additionally, we will explain the main parameters in the next subsection.

4.1 Parameter Settings

The algorithm is unsupervised in the sense that the objects orientation in space as well as its hierarchical structure are unknown. However, prior knowledge is introduced by the parameter settings, which is why we want to explain these in a bit more detail. The number of parts to be composed E can be set freely, but it has significant implications for the structure of the model. E.g. the number of layers will be much lower for higher values of E . Moreover, for different objects and also for different parts inside the same object the optimal value of E can be different. Thus it may need further research to determine this parameter automatically. We set $E = 2$ in the bottom-up as well as the top-down process, as this is the minimal number. The image sizes during the experiments are fixed to have a diagonal of 250 pixels, while the Gabor filters are quadratic

with a size of 17 pixels. The Gabor filter is rotated in 180 degrees in 10 degree steps. All higher layer parts are searched over the full rotation of 360 degrees in 10 degree steps. The active parameters Δ_L can be set for each layer individually. During learning we set $\delta_\alpha = 0$ and $\delta_x = 1$ for all layers. As the higher layers work directly on the response maps of the lower layers, the higher layers become more and more invariant to positional transformation. We start with a part size of 35 pixels and double the size for each layer. However, when the parts reach the border of the template, it automatically increases up to a certain amount. The threshold for suppression the hard inhibition of overlapping filters is set to $\tau = .05$. With this value, the filters in the part templates are well spaced. The overlap becomes more important in the higher layers as the training data becomes more scarce. There we gradually increase tolerance to overlap. Experimental results have shown that during the EM-type learning, a part is converged to stable solution after around 5 iterations, therefore we fix the number of iterations per part to this value.

Patch sampling strategies. An important step during the compositional learning process is the proposal of an initial part model (Step 9&10 in Algorithm 3). A common approach is to sample positions in the image randomly, followed by cutting out the response maps at those positions and subsequently learning an initial model from these patches. However, as the parts become more complex the part responses are more localized and it becomes less likely to sample informative patches. In order to overcome this issue, we sample the positions according to the observed distribution of responses instead of a uniform distribution. We found that this resolves the issue satisfactory.

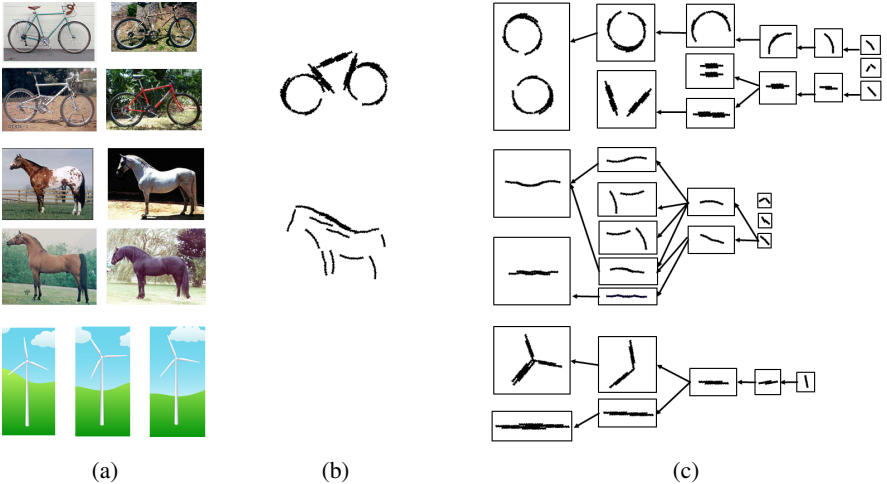


Fig. 4: Learned hierarchical compositional models. (a) Samples from the training data. (b) The object representation after the top-down process. (c) The compositional structure learned with the bottom-up process.

4.2 Unsupervised Learning of Hierarchical Compositional Models

In Figure 4 we illustrate hierarchical compositional models that were learned from different images with the proposed approach. Figure 4b shows that the proposed learning scheme is able to learn meaningful object representations from natural images. The bottom-up process is able to automatically induce a compositional structure with variable number of parts and layers (see Figure 4c). However, for articulated objects such as the wind turbine the process can not couple the top layer parts as we currently assume that $\delta_\alpha = 0$ and thus that the relative orientation between parts stays the same. Another important constraint is that the parts are only allowed to compose with other parts from the same layer. As a consequence of the fixed number of part compositions E this leads to suboptimal models and ultimately to the top layer of the bottom up procedure not resembling the full object shape (see horse model). However, the top-down process can robustly recover the object shape as far as the learned parts allow it.

4.3 Domain Adaptation

Domain adaptation is the task of transferring knowledge between different data distributions. The four domain dataset was proposed in [26] based on the three domain dataset published in [27]. The four-domain dataset is composed of 10 classes that are commonly shared across the following datasets: *Amazon* with images downloaded from Amazon; *DSLR* with high-resolution images; *Webcam* with low-resolution images and images from *Caltech256* [28]. In each domain the resolution, lighting conditions, background, the object textures and positions in space vary significantly. Therefore the images are considered as coming from different domains. For the experiments, we follow the standard evaluation protocol [26]. We test two evaluation setups. In the semi-supervised setting, the algorithm also has access to a small amount of data from the target domain, whereas in the unsupervised setting the training images are only sampled from the source domain. As proposed in [3] we learn the compositional model from the training images and feed the response maps of all parts from the model to the spatial pyramid matching method [29]. This method partitions the maps into 1, 4, and 16 image regions and aggregates the maximal responses in these regions into a feature vector. A Multi-Class SVM is then trained based on these features. In Table 1 we compare our results with several other approaches on the unsupervised learning task.

Our approach achieves the best results in three out of eight classification experiments. Moreover, we tested the HABM approach from [3] on the same data and in many cases outperformed this approach by a wide margin. Interestingly, both HABM and our approach work with exactly the same Gabor filters but in a completely different structural relation. This result underlines the benefits of an automatic structure induction process. Clearly, different objects have different hierarchical structures and the experimental findings suggest that exploiting this property improves the classification accuracies. For the semi-supervised learning setup, we have compared our classification performance to several recent approaches. The results are listed in Table 2. Again we outperform the HABM approach significantly in the experiments. In two out of eight experiments our result is better than the current state-of-the-art. However, we point out that in contrast to

Methods	$C \rightarrow A$	$C \rightarrow D$	$A \rightarrow C$	$A \rightarrow W$	$W \rightarrow C$	$W \rightarrow A$	$D \rightarrow A$	$D \rightarrow W$
KSVD [30]	20.5 ± 0.8	19.8 ± 1.0	20.2 ± 0.9	16.9 ± 1.0	13.2 ± 0.6	14.2 ± 0.7	14.3 ± 0.3	46.8 ± 0.8
SGF [31]	36.8 ± 0.5	32.6 ± 0.7	35.3 ± 0.5	31.0 ± 0.7	21.7 ± 0.4	27.5 ± 0.5	32.0 ± 0.4	66.0 ± 0.5
GFK [26]	40.4 ± 0.7	41.1 ± 1.3	37.9 ± 0.4	35.7 ± 0.9	29.3 ± 0.4	35.5 ± 0.7	36.1 ± 0.4	79.1 ± 0.7
SIDL [32]	45.4 ± 0.3	42.3 ± 0.4	40.4 ± 0.5	37.9 ± 0.9	36.3 ± 0.3	38.3 ± 0.3	39.1 ± 0.5	86.2 ± 1.0
HABM [3]	53.7 ± 4.7	43.2 ± 4.9	41.2 ± 1.6	28.1 ± 2.0	25.8 ± 1.6	33.5 ± 2.9	34.6 ± 3.7	68.2 ± 2.9
OURS	62.3 ± 3.4	43.7 ± 2.9	54.0 ± 2.4	33.3 ± 1.7	29.5 ± 1.1	35.0 ± 3.6	33.1 ± 2.4	65.6 ± 3.8

Table 1: **Unsupervised** learning: Classification scores on the four domain benchmark.

Methods	$C \rightarrow A$	$C \rightarrow D$	$A \rightarrow C$	$A \rightarrow W$	$W \rightarrow C$	$W \rightarrow A$	$D \rightarrow A$	$D \rightarrow W$
Metric [27]	33.7 ± 0.8	35.0 ± 1.1	27.3 ± 0.7	36.0 ± 1.0	21.7 ± 0.5	32.3 ± 0.8	30.3 ± 0.8	55.6 ± 0.7
SGF [31]	40.2 ± 0.7	36.6 ± 0.8	37.7 ± 0.5	37.9 ± 0.7	29.2 ± 0.7	38.2 ± 0.6	39.2 ± 0.7	69.5 ± 0.9
GFK [26]	46.1 ± 0.6	55.0 ± 0.9	39.6 ± 0.4	56.9 ± 1.0	32.8 ± 0.1	46.2 ± 0.6	46.2 ± 0.6	80.2 ± 0.4
FDDL [33]	39.3 ± 2.9	55.0 ± 2.8	24.3 ± 2.2	50.4 ± 3.5	22.9 ± 2.6	41.1 ± 2.6	36.7 ± 2.5	65.9 ± 4.9
Landmark [34]	56.7	57.3	45.5	46.1	35.4	40.2	-	-
HMP [35]	67.7 ± 2.3	70.2 ± 5.1	51.7 ± 4.3	70.0 ± 4.2	46.8 ± 2.1	61.5 ± 3.8	64.7 ± 2	76.0 ± 4
SDDL [36]	49.5 ± 2.6	76.7 ± 3.9	27.4 ± 2.4	72.0 ± 4.8	29.7 ± 1.9	49.4 ± 2.1	48.9 ± 3.8	72.6 ± 2.1
SIDL [32]	$50.0 \pm .5$	57.1 ± 0.4	41.5 ± 0.8	57.8 ± 0.5	40.6 ± 0.4	51.5 ± 0.6	± 0.6	80.2 ± 0.4
DASH-N [37]	71 ± 2.2	81.4 ± 3.5	54.9 ± 1.8	75.5 ± 4.2	50.2 ± 3.3	70.4 ± 3.2	68.9 ± 2.9	77.1 ± 2.8
HABM [3]	68.3 ± 2.3	57.4 ± 6.0	52.7 ± 3.0	54.8 ± 2.8	42.2 ± 3.1	57.1 ± 3.5	60.1 ± 3.2	79.7 ± 2.5
OURS	72.2 ± 0.7	58.1 ± 5.1	58.5 ± 1.2	53.4 ± 1.2	47.6 ± 1.8	61.7 ± 3.2	65.6 ± 2.8	78.5 ± 2.0

Table 2: **Semi-supervised** learning: Classification scores on the four domain benchmark.

most other competitive approaches, our approach was not primarily designed as a feature extractor for a classifier. Instead it is a generative, fully probabilistic object model that can be applied to a wide range of applications.

5 Conclusion & Future Work

We propose a greedy EM-type algorithm for unsupervised learning of hierarchical compositional models. The learning procedure is simple and makes it possible to learn a hierarchical compositional model that is probabilistic and fully generative up to the pixel level. The greedy algorithm can be universally applied for a bottom-up induction of the model structure as well as within a top-down model building process. Experimental results show that the learned hierarchical compositional representation achieves competitive results for a domain adaptation task compared to feature based approaches and hierarchical models with manually designed structure.

We see many future directions for research. The iterative application of the bottom-up learning and the top-down correction process might increase the robustness to clutter and changes in the lighting conditions. Furthermore, the Learning of the spatial distribution between parts, instead of a pre-specifying their expected relative movement, would make the modeling of articulated objects possible. Finally, we would propose to extend the model to jointly represent multiple objects, thus allowing for part sharing and the assessment of semantic relations between objects.

References

1. Fidler, S., Boben, M., Leonardis, A.: Learning a hierarchical compositional shape vocabulary for multi-class object representation. arXiv preprint arXiv:1408.5516 (2014)
2. Zhu, L.L., Lin, C., Huang, H., Chen, Y., Yuille, A.: Unsupervised structure learning: Hierarchical recursive composition, suspicious coincidence and competitive exclusion. In: Computer vision—eccv 2008. Springer (2008) 759–773
3. Dai, J., Hong, Y., Hu, W., Zhu, S.C., Wu, Y.N.: Unsupervised learning of dictionaries of hierarchical compositional models. In: Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on, IEEE (2014) 2505–2512
4. Wong, A., Yuille, A.: One shot learning via compositions of meaningful patches. In: ICCV. (2015)
5. Fidler, S., Leonardis, A.: Towards scalable representations of object categories: Learning a hierarchy of parts. In: Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, IEEE (2007) 1–8
6. Si, Z., Gong, H., Zhu, S.C., Wu, Y.N., et al.: Learning active basis models by em-type algorithms. *Statistical Science* **25**(4) (2010) 458–475
7. Wu, Y.N., Si, Z., Gong, H., Zhu, S.C.: Learning active basis model for object detection and recognition. *International journal of computer vision* **90**(2) (2010) 198–235
8. Hong, Y., Si, Z., Hu, W., Zhu, S.C., Wu, Y.N.: Unsupervised learning of compositional sparse code for natural image representation. *Quarterly of Applied Mathematics* **72** (2013) 373–406
9. Si, Z., Zhu, S.C.: Learning and-or templates for object recognition and detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **35**(9) (2013) 2189–2205
10. Fukushima, K.: Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics* **36**(4) (1980) 193–202
11. Felzenszwalb, P.F., Huttenlocher, D.P.: Pictorial structures for object recognition. *International Journal of Computer Vision* **61**(1) (2005) 55–79
12. Epshtein, B., Uliman, S.: Feature hierarchies for object classification. In: Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on. Volume 1., IEEE (2005) 220–227
13. Jin, Y., Geman, S.: Context and hierarchy in a probabilistic image model. In: Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on. Volume 2., IEEE (2006) 2145–2152
14. Hinton, G.E., Osindero, S., Teh, Y.W.: A fast learning algorithm for deep belief nets. *Neural computation* **18**(7) (2006) 1527–1554
15. Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., Poggio, T.: Robust object recognition with cortex-like mechanisms. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **29**(3) (2007) 411–426
16. Ommer, B., Buhmann, J.M.: Learning the compositional nature of visual object categories for recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **32**(3) (2010) 501–516
17. Todorovic, S., Ahuja, N.: Unsupervised category modeling, recognition, and segmentation in images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **30**(12) (2008) 2158–2174
18. Zhu, L., Chen, Y., Yuille, A.: Learning a hierarchical deformable template for rapid deformable object parsing. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **32**(6) (2010) 1029–1043

19. Geman, S., Potter, D.F., Chi, Z.: Composition systems. *Quarterly of Applied Mathematics* **60**(4) (2002) 707–736
20. Yuille, A.L., Hallinan, P.W., Cohen, D.S.: Feature extraction from faces using deformable templates. *International journal of computer vision* **8**(2) (1992) 99–111
21. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. *IEEE Transactions on Pattern Analysis & Machine Intelligence* (6) (2001) 681–685
22. Olshausen, B.A., et al.: Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**(6583) (1996) 607–609
23. Riesenhuber, M., Poggio, T.: Hierarchical models of object recognition in cortex. *Nature neuroscience* **2**(11) (1999) 1019–1025
24. Mallat, S.G., Zhang, Z.: Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on* **41**(12) (1993) 3397–3415
25. Pietra, S.D., Pietra, V.D., Lafferty, J.: Inducing features of random fields. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **19**(4) (1997) 380–393
26. Gong, B., Shi, Y., Sha, F., Grauman, K.: Geodesic flow kernel for unsupervised domain adaptation. In: *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, IEEE (2012) 2066–2073
27. Saenko, K., Kulis, B., Fritz, M., Darrell, T.: Adapting visual category models to new domains. In: *Computer Vision—ECCV 2010*. Springer (2010) 213–226
28. Griffin, G., Holub, A., Perona, P.: Caltech-256 object category dataset. (2007)
29. Lazebnik, S., Schmid, C., Ponce, J.: Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In: *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*. Volume 2., IEEE (2006) 2169–2178
30. Aharon, M., Elad, M., Bruckstein, A.: K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *Signal Processing, IEEE Transactions on* **54**(11) (2006) 4311–4322
31. Gopalan, R., Li, R., Chellappa, R.: Domain adaptation for object recognition: An unsupervised approach. In: *ICCV*. (2011)
32. Ni, J., Qiu, Q., Chellappa, R.: Subspace interpolation via dictionary learning for unsupervised domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2013) 692–699
33. Yang, M., Zhang, L., Feng, X., Zhang, D.: Fisher discrimination dictionary learning for sparse representation. In: *ICCV*. (2011)
34. Gong, B., Grauman, K., Sha, F.: Connecting the dots with landmarks: Discriminatively learning domain-invariant features for unsupervised domain adaptation. In: *ICML*. (2013)
35. Bo, L., Ren, X., Fox, D.: Hierarchical matching pursuit for image classification: Architecture and fast algorithms. In: *Advances in neural information processing systems*. (2011) 2115–2123
36. Shekhar, S., Patel, V., Nguyen, H., Chellappa, R.: Generalized domain-adaptive dictionaries. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2013) 361–368
37. Nguyen, H.V.: Non-linear and sparse representations for multi-modal recognition. (2013)